

On Measuring Bias in Online Information

Evaggelia Pitoura, Panayiotis Tsaparas
Department of Computer Science and
Engineering
University of Ioannina
{pitoura,tsap}@cs.uoi.gr

Serge Abiteboul
INRIA & ENS Cachan
serge.abiteboul@inria.fr

Giorgos Flouris, Irini Fundulaki,
Panagiotis Papadakos
Institute of Computer Science, FORTH
{fgeo,fundul,papadako}@ics.forth.gr

Gerhard Weikum
Max Planck Institute for Informatics, Germany
weikum@mpi-inf.mpg.de

ABSTRACT

Bias in online information has recently become a pressing issue, with search engines, social networks and recommendation services being accused of exhibiting some form of bias. In this vision paper, we make the case for a systematic approach towards measuring bias. To this end, we discuss formal measures for quantifying the various types of bias, we outline the system components necessary for realizing them, and we highlight the related research challenges and open problems.

1. INTRODUCTION

We live in an information age where the majority of our diverse information needs are satisfied online by search engines, social networks and media, news aggregators, e-shops, vertical portals, and other online information providers (**OIPs**). For every request we submit to these providers, a combination of sophisticated algorithms produce a ranked list of the most relevant results tailored to our profile. These results play an important role in guiding our decisions (e.g., where should I dine, what should I buy, which jobs should I apply to), in shaping our opinions (e.g., who should I vote for), and in general in our view of the world.

Undoubtedly, the various OIPs help us in managing and exploiting the abundance of available information. But, at the same time, the convenient and effective way in which the OIPs satisfy our information needs has limited our information seeking abilities, and has rendered us overly dependent on them. We rarely wonder whether the returned results are properly representative of all different viewpoints, and we seldom escape the echo chambers and filter bubbles created by personalization. We have come to accept such results as the “de facto” truth.

There are increasingly frequent reports of OIPs ex-

hibiting some form of bias. For instance, in the recent US presidential elections, Google was accused of being biased against Donald Trump¹ and Facebook of contributing to the post-truth politics². Google search has been accused of being sexist or racist when returning images for queries such as “nurse” or “hair-styling”³, and prejudiced when answering queries about holocaust⁴. Similar accusations have been made for Flickr, Airbnb and LinkedIn. In fact, the problem of understanding and addressing bias is considered a high-priority problem for machine learning algorithms and AI for the next few years⁵.

The problem has attracted some attention in the data management community as well [33]. In this paper, we make the case for a systematic approach to addressing the problem of bias in the data provided by the OIPs. Addressing bias involves many steps. Here, we focus on the very first step, that, of defining and measuring bias.

According to the Oxford English Dictionary⁶, bias is “an inclination or prejudice for or against one person or group, especially in a way considered to be unfair”, and as “a concentration on or interest in one particular area or subject”. When it comes to bias in OIPs, we make the distinction between *subject bias* and *object bias*. Subject bias refers to bias towards the users that receive a result. Subject bias appears when different users receive different content based on user attributes that should be protected, such as gender, race, ethnicity, or religion. Object bias refers to biases in the content of the results, for example, when an aspect is disproportionately represented in a result.

In the remaining of this paper, we propose a number of formal measures for both subject and object bias. Then, we present the basic components required by a

¹<https://www.theguardian.com/us-news/2016/sep/29/donald-trump-attacks-biased-lester-holt-and-accuses-google-of-conspiracy>

²<https://www.theguardian.com/us-news/2016/nov/16/facebook-bias-bubble-us-election-conservative-liberal-news-feed>

³<http://fusion.net/story/117604/looking-for-ceo-doctor-cop-in-google-image-search-delivers-crazy-sexist-results/>

⁴<http://www.bbc.com/news/technology-38379453>

⁵<https://futureoflife.org/ai-principles/>

⁶<https://en.oxforddictionaries.com/definition/bias>.

system for realizing these measures towards identifying bias in an OIP. Finally, we provide a synopsis of the research challenges and open problems in defining and measuring bias in online information.

2. RELATED WORK

In the field of machine learning, there is an increasing concern about the potential risks of data-driven approaches in decision making algorithms [2, 3, 17, 24, 30, 33], raising a call for equal opportunities by design [21]. Biases can be introduced at different stages of the design, implementation, training and deployment of machine learning algorithms. There are reports for discriminatory ads based on either race [32, 34], or gender [11], and recommendation algorithms may show different prices to different users [19]. AdFisher [9] runs browser-based experiments to explore how user behaviors and profiles affect ads and if they can lead to seemingly discriminatory ads. Consequently, there are efforts for defining principles of accountable algorithms⁷, for auditing algorithms by detecting discrimination [31] and for debiasing rating approaches [1]. There is a special interest for racial fairness and fair classifiers [20, 39, 40, 7], in order to make sure that groups receive ads based on population proportions [11] and reduce the discrimination degree of algorithms against individuals of a protected group [15]. Other efforts try to ensure temporal transparency for policy changing events in decision making systems [14]. Finally, tools that remove discriminating information⁸, help in understanding opposing opinions⁹, flag fake news¹⁰, increase transparency of personalization algorithms¹¹, or show political biases of Facebook friends and news feed¹² have recently appeared.

Another branch of research focuses on how bias can affect users. According to field studies, users of search engines trust more the top-ranked search results [28] and biased search algorithms could shift the voting preferences of undecided voters by as much as 20% [12]. Since most users try to access information that they agree with [22], the personalization and filtering algorithms used by search engines lead to echo chambers and filter bubbles that reinforce bias [4, 18]. This is also evident in social media where platforms strengthen users existing biases [25], minimizing the exposure to different opinions [36]. Rating bubbles emerge especially when positive social influence accumulates, while crowd correction neutralizes negative influence [27].

Previous studies have looked at individual aspects of bias, such as geographical (i.e. whether sites from certain countries are covered more) [35], or temporal (recommending recent and breaking news) [6]. Other ap-

proaches try to examine how bias can be measured [26] and if search engines can partially mitigate the richer nature of the Web and give new sites an increased chance of being discovered [16]. A number of studies have explored the presence of bias in media sources through human annotations [5], by exploiting affiliations [37] and the impartiality of messages [38], while [23] tries to quantify bias in twitter data. There is clearly a need for a systematic way to measure bias, and in this vision paper we set the groundwork for a complete system for this task.

3. TYPES OF BIAS

We consider bias in the context of *topics*. In particular, we would like to test whether an OIP is biased with regards to a given topic. A topic may be a very general or a very specific one down to the granularity of a single search query. For example, we may want to test whether an OIP provides biased results for news topics such as “Brexit” and “US Elections”, for people such as “Donald Trump”, for general issues such as “abortion” and “gun control”, for transactional queries such as “air tickets”, “best burger”, or even topics such as “famous people”.

We distinguish between two types of bias, namely *subject* and *object* bias. Subject bias refers to bias against the users receiving the information, while object bias looks at bias in the actual content.

For subject bias, we assume that some of the attributes that characterize the user of an OIP are *protected attributes*, for example, race, or gender. We consider that there is subject bias, if the values of these attributes influence the results presented to the users. This type of bias can be encountered for example in the case of a query about jobs, when women get results about lowered paid jobs compared to the results shown to men. Note that subject bias can also appear as a consequence of some hidden dependencies between the protected and unprotected attributes, even when the values of the protected attributes are not used directly in computing the results (e.g., see [13]). For example, the home location of users may imply their race.

Object bias refers to the requirement that content is unbiased. This type of bias may appear even when we have just one user. For example, an instance of this kind of bias occurs when an OIP promotes its own services over the competitive ones, or, when the results for queries about a political figure take an unjustifiable favorable, or non-favorable position towards this politician (independently of the user receiving the results).

In analogy to protected attributes for users, we assume *differentiating aspects* for topics. For instance, for a controversial topic such as “abortion” or “gun control”, the differentiating aspect would be the stance (pro, or against). For a topic such as “famous people”, we may want to test whether the results are biased towards men over women, or, favor people from specific countries, or, over-represent, say, artists over scientists. Finally, for a topic such as “US Elections”, where we want to test if there is bias towards some party, the dif-

⁷<http://www.fatml.org/resources/principles-for-accountable-algorithms>

⁸<http://www.debiasyourself.org/>

⁹<https://www.escapeyourbubble.com/>

¹⁰<http://www.theverge.com/2016/12/15/13960062/facebook-fact-check-partnerships-fake-news>

¹¹<https://facebook.tracking.exposed/>

¹²<http://politecho.org/>

differentiating aspect would be the party (“Democrats” or “Republicans”).

In a sense, addressing subject bias can be regarded as a counterweight to machine-learning and personalization algorithms that try to differentiate the needs of various user groups, so that the algorithms do not discriminate over specific protected attributes. On the other hand, addressing object bias has some similarity to result diversification [10]. However, diversity is more related to coverage, in the sense that we want all different aspects of a topic to appear in the result, even the rarest ones. In terms of object bias, we want the aspects to be represented proportionally to a specific “ground truth”.

A commonly encountered case is the case of a combined subject-object bias appearing when a specific facet is over-represented in the results presented to specific user populations, e.g., democrats get to see more pro-Clinton articles than republicans. This type of bias is also related to *echo chambers*, i.e., the situation in which information, ideas, or beliefs are amplified, exaggerated or reinforced inside groups of equally-minded people. Since similar people may be interested in specific aspects of a topic, as a result the content they create, consume, or prefer is biased towards these aspects. Then, the information presented to them may reflect this bias and by doing so possibly amplify the bias, creating a bias-reinforcement cycle. In such cases, there is often some relation between the protected attributes of the users and the differentiating aspects of the topic.

Finally, note that, although there is a negative connotation associated with the term bias, some forms of “positive bias” may also be introduced as a result of legislations, and governmental or, other policies, for example to foster gender equality, enforce diversity, or, to protect the public interest.

4. BIAS MEASURES

We will now introduce definitions of measures of bias, for both object and subject bias. Our measures are to some extent influenced by models of fairness in classification systems (e.g., [11]). Our goal is not to be overly formal, but instead we provide such definitions as a means to make the research challenges involved more concrete.

To simplify the discussion, for the following, we assume that the topic for which we want to measure bias is just a single query q . We can generalize the definitions to a set of queries by adopting some aggregation measure of the metrics for a single query.

4.1 Subject Bias

Let U be the user population. Intuitively, we would like the information provided to the users to not be influenced by their protected attributes. For simplicity assume a binary protected attribute that divides users into a protected class P and an unprotected class \bar{P} . For example, if the protected attribute is gender, P may denote the set of men and \bar{P} the set of women.

We use R_u to denote the ranked list of results for

user u for query q , and $R = \cup_{u \in U} R_u$ to denote the union of all results for all users in U . A result $r \in R$ may be a single document, or a class of documents that we consider to be “equivalent” (e.g., a set of articles that are all in favor of a political party, or a collection of job postings that are all high-paid).

Let $Pr(u, r)$ denote the probability that a user u receives result r , where the probability is defined over the result lists of all users in U . We propose the following definition.

DEFINITION 1 (SUBJECT BIAS-1). *An online information provider is subject-unbiased if for every possible result $r \in R$ for query q , it holds:*

$$|Pr(u, r|u \in P) - Pr(u, r|u \notin P)| \leq \epsilon$$

for some small $\epsilon \geq 0$.

Intuitively, the definition says that the probability that a user sees some result does not depend on the membership of the user to the protected class. There are different ways to compute the probability $Pr(u, r)$. For example, we can define it as the fraction of users that have r in their top- k results, or even take into account the exact position in which the result r appears in the ranking.

We also consider a distance-based definition. This definition is inspired by [11] and assumes that we can define an appropriate distance measure D_u between the users in U . D_u must be such that it does not consider the protected attributes of the users (or, any information inferred by them). We also assume that given the ranked lists of results R_{u_1} and R_{u_2} for users u_1 and u_2 respectively there is a distance function $D_r(R_{u_1}, R_{u_2})$ that measures the distance between the rankings. D_r can be defined by employing existing distance metrics between ranked lists, or using a geometric embedding of the ranked lists that takes the document equivalence, and the importance of position into account. We propose the following alternative definition of subject bias.

DEFINITION 2 (SUBJECT BIAS-2). *An online information provider is subject-unbiased if for any pair of users u_1 and u_2 , it holds:*

$$D_r(R_{u_1}, R_{u_2}) \leq D_u(u_1, u_2)$$

The intuition behind this definition is that similar users should receive similar result lists.

4.2 Object Bias

For now, assume just one user u , and a single query q . Let A be the differentiating aspect, and let $\{a_1, \dots, a_m\}$ be the values of A . For example, in the case that the query is about elections, a_1, \dots, a_m would correspond to the different parties that participate in the elections. We also assume that each result is annotated with the values of aspect A , meaning that the result is about these values. It is possible that a result is annotated with multiple values, or with none of the values.

We provide again two definitions for object bias. As in the case of subject bias, let $Pr(u, a_i)$ be the probability that user u receives a result annotated with value

a_i (e.g., one possible definition is this to be defined as the fraction of the top- k results that are about a_i). Furthermore, we also assume we have the “ground truth” which is in the form of probabilities $Pr_T(a_i)$ for all the aspect values which captures the relative popularity of each aspect value (e.g., the support of a party as measured by polls). We propose the following definition of object bias.

DEFINITION 3 (OBJECT BIAS-1). *An online information provider is object-unbiased with respect to aspect A if for every value a_i of A , it holds:*

$$|Pr(u, a_i) - Pr_T(a_i)| \leq \epsilon$$

for some small $\epsilon \geq 0$.

For the distance based definition, we assume that we have ground truth in the form of a ranking R_T which captures the ideal unbiased ranking. We propose the following definition.

DEFINITION 4 (OBJECT BIAS-2). *An online information provider is object-unbiased if it holds:*

$$D_r(R_u, R_T) \leq \epsilon$$

Note that the differentiating aspects do not appear explicitly in Definition 4. Such aspects will be encapsulated in the definition of D_r .

5. A SYSTEM FOR MEASURING BIAS

We now look at the various challenges involved in realizing these definitions for measuring the bias of an OIP. The OIP may be a search engine, a recommendation service, the search or news feed service of a social network. We will treat the OIP as a black-box and assume that we can access it only through the interface that the OIP provides, e.g., through search queries.

In Figure 1, we present the main components that a system for measuring bias should provide. We call this system **BIASMETER**.

BIASMETER takes as input: (1) the user population U for which we want to measure the (subject) bias, (2) the set P of the protected attributes of U , (3) the topic T for which we want to test the (object) bias and (4) the set A of the differentiating aspects of the topic T .

For simplicity, we assume that the protected and the differentiating aspects (P and A , respectively) are given as input. A more daunting task would be to infer these attributes.

Next, we describe in some detail the components of **BIASMETER**. Given the topic T and the differentiating aspects A , the goal of the *query generator* is to produce an appropriate set of queries to be submitted to the OIP under consideration. For example, if the OIP is a search engine, the query generator should produce a set of search queries. For instance, to test about the topic “US elections”, the generator may produce a variety of queries, including queries referring to specific political parties. To produce queries that best represent the topic and the aspects, the query-generator may need to use a related knowledge base.

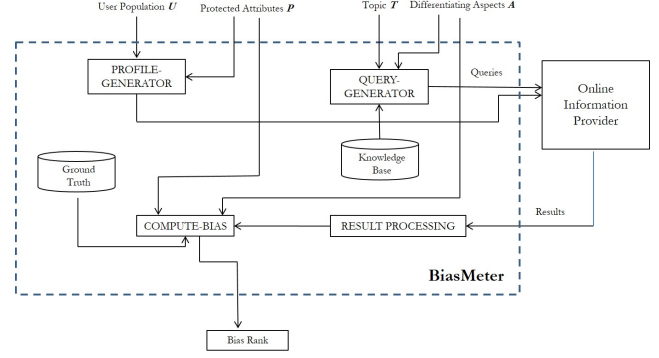


Figure 1: System components.

The *profile generator* takes as input the user population U and the set of protected attributes P and produces as output a set of user profiles appropriate for testing whether the OIP discriminates over users in U based on the protected attributes in P . For example, if we want to test whether the provider discriminates based on the gender of the user that issues the query, we may need to generate an appropriate number of representative accounts of users of different gender.

There are many issues of both a theoretical and a practical nature regarding generating profiles. For example, we must ensure that the generated profiles are an appropriate sample of U that represents all values of the protected attributes. Furthermore, we should ensure that the characteristics of the users in the sample are the same with respect to all other attributes, so as to avoid the effect of confounding factors. This raises issues similar to those met when selecting people for opinion polls, surveys, etc. From a more practical view, we need to assemble users with the specific profiles and ask them to issue the queries (for example using a crowd-sourcing platform, such as Mechanical Turk), or generate artificial accounts of such users. Doing that at a large scale is a challenging problem.

The *result processing* component takes as input the results from the OIP and applies machine learning and data mining algorithms such as topic modeling and opinion mining to determine the value of the differentiating aspects. For example, if the topic is “gun control”, we need to determine whether a specific result takes a positive, neutral or negative stand.

A module central to a system for measuring bias is the *ground-truth* module. Obtaining the ground truth is hard in practice. One possible way to overcome this problem is to consider a collection of OIPs, and require that the results of the OIP in question are close on average to the results of all of them. The idea in this case is that by looking at multiple providers we capture the wisdom of the crowd, to which the provider we test should be close. Other approaches include: (a) looking at the distribution in the general population, and (b) using external sources, e.g., Wikipedia.

Finally, the *compute-bias* component calculates the bias of the OIP, using our bias metrics.

Examples

We now describe some (relatively) concrete examples of how we envision the operation of the different system components.

Subject bias. We begin with an example of subject bias. Consider the case that a computer science graduate is looking for an engineering job. We want to identify if there is bias with respect to gender: whether women are shown different jobs than men. In this case, the OIP is a search engine (possibly a vertical search engine that specializes in job search). The user population U is all CS graduates, (we could further limit the population by considering a specific location, or specific school), the protected attribute P is gender, and the topic T is “engineering job”.

Given the topic, the query generator should produce queries to the search engine that capture the information need of the user; in this case, queries of the form “computer engineering jobs”, “CS employment opportunities”, etc. The profile generator should output an appropriate sample of users. We need samples of men and women, that have very similar characteristics with respect to other attributes such as grades, skills, background, ethnicity, etc, to avoid differences that may appear due to attribute correlations. Given the set of queries and the population sample, we pose the queries to the search engine and obtain the results.

Using the metrics discussed in Section 4, depending on the threshold we set, we can determine if there is bias in the OIP with respect to the protected attribute and the specific topic. Note that the cause of bias is not specified in our result; we just know that the search engine treats women and men differently.

Object bias. We will now consider an example of object bias. Consider the case that a user is looking for information on US elections. We want to determine if the OIP is biased towards one of the political parties.

We assume again that the OIP is a search engine. The user population U is all users in US, although we could consider a specific region or demographic population. There is no protected attribute P in this case since we are measuring object bias. The topic T is “US elections” and the differentiating aspect A is the political party.

Similar to before, the query generator should output a collection of queries that capture the information need of the users for the topic. The profile generator in this case should output a random sample of users. Different users may see different results depending on personalization parameters, but with a large enough sample of users, we expect these differences to be “averaged out”. The result processing stage will label the results according to the political party that they favor (if any). We can now compute the metrics we described in Section 4. To obtain the ground truth popularity of the different parties, we can use external sources, such as polls, that measure the party popularity and intention to vote.

6. RESEARCH CHALLENGES

Measuring the bias of an OIP raises many research challenges. Here we highlight the most important ones.

Defining a bias measure. Bias is multifaceted. Here, we abstracted the many forms of bias, through the notions of differentiating aspects and protected attributes. However, there are often correlations among the attributes and the aspects making it difficult to single them out. Quantifying bias is also complex. Our definitions rely on probabilities or the existence of appropriate distance metrics between users and rankings. Estimating such probabilities and defining appropriate distances remains an open issue.

Obtaining the ground truth. Acquiring the ground truth is one of the most formidable tasks in measuring bias. One possible approach is to employ manual effort for constructing or evaluating the results, in a similar way that search engines evaluate the ranking quality. However, this comes against cognitive user bias such as confirmation bias where users prefer information that fits their own beliefs. Thus, it may be hard to even find objective evaluators, or users that would accept our measurements of bias as objective ones.

Obtaining the user population. In order to measure bias with respect to a protected attribute P (e.g. gender), we need to generate large samples of user accounts for the different values of P (e.g., women and men). We should ensure that the distribution of all other characteristics is near identical for all other attributes. Careful statistical analysis is required to ensure the statistical significance of our results.

Engineering and technical challenges. The query generation and result processing components involve a variety of data mining and machine learning algorithms for identifying keywords that describe an information need, or understanding the topic and stance of a specific document. Towards this end we need modules for knowledge representation, record linkage, entity detection and entity resolution, sentiment detection, topic modeling, and more.

Legal issues. The bias detection system we described could be significantly simplified if we had access to the internal workings of the OIP (e.g., for sampling users with specific demographics, or obtaining query results with no personalization). Clearly, this is impossible for an entity outside the OIP. However, it may be in the interest of the governments to create legislation that obliges OIPs to provide access to sufficient data for measuring bias. There is a growing literature advocating the systematic auditing of algorithms [8, 29, 31].

7. CONCLUSIONS

In this vision paper, we argue about the importance of a systematic approach for measuring the bias of the information we get from online information providers. As more and more people rely on online sources to get informed and make decisions, such an approach is of central value. Building a system for measuring bias raises many research challenges, some of which we have highlighted in this paper.

Measuring bias is just the first step. To counteract bias, many more steps are needed including identifying the sources of bias and taking actions to address it.

8. REFERENCES

- [1] G. Adomavicius, J. Bockstedt, C. Shawn, and J. Zhang. *De-biasing user preference ratings in recommender systems*, volume 1253, pages 2–9. CEUR-WS, 2014.
- [2] S. Barocas and A. D. Selbst. Big Data’s Disparate Impact. *SSRN eLibrary*, 2014.
- [3] T. Bolukbasi, K. Chang, J. Y. Zou, V. Saligrama, and A. Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *CoRR*, abs/1607.06520, 2016.
- [4] E. Bozdag. Bias in algorithmic filtering and personalization. *Ethics and Inf. Technol.*, 15(3):209–227, Sept. 2013.
- [5] C. Budak, S. Goel, and J. M. Rao. Fair and balanced? quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly*, 80(S1):250–271, 2016.
- [6] A. Chakraborty, S. Ghosh, N. Ganguly, and K. P. Gummadi. Can trending news stories create coverage bias? on the impact of high content churn in online news media. In *Computation and Journalism Symposium*, 2015.
- [7] S. Corbett-Davies, E. Pierson, A. Feller, S. Goel, and A. Huq. Algorithmic decision making and the cost of fairness. *CoRR*, abs/1701.08230, 2017.
- [8] K. Crawford. Can an algorithm be agonistic? ten scenes from life in calculated publics. *Science, Technology, & Human Values*, 41(1):77–92, 2016.
- [9] A. Datta, M. C. Tschantz, and A. Datta. Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies*, 2015(1):92–112, 2015.
- [10] M. Drosou and E. Pitoura. Search result diversification. *SIGMOD Record*, 39(1):41–47, 2010.
- [11] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. S. Zemel. Fairness through awareness. In *ITCS*, pages 214–226, 2012.
- [12] R. Epstein and R. E. Robertson. The search engine manipulation effect (seme) and its possible impact on the outcomes of elections. *PNAS*, 112(20), 2015.
- [13] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian. Certifying and removing disparate impact. In *KDD*, pages 259–268, 2015.
- [14] M. Ferreira, M. B. Zafar, and K. P. Gummadi. The case for temporal transparency: Detecting policy change events in black-box decision making systems. *arXiv preprint arXiv:1610.10064*, 2016.
- [15] B. Fish, J. Kun, and Á. D. Lelkes. A confidence-based approach for balancing fairness and accuracy. In *SDM*, pages 144–152, 2016.
- [16] S. Fortunato, A. Flammini, F. Menczer, and A. Vespignani. Topical interests and the mitigation of search engine bias. *Proceedings of the National Academy of Sciences*, 103(34):12684–12689, 2006.
- [17] S. Hajian, F. Bonchi, and C. Castillo. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *KDD*, pages 2125–2126. ACM, 2016.
- [18] A. Hannak, P. Sapiezynski, A. Molavi Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson. Measuring personalization of web search. In *WWW*, pages 527–538. ACM, 2013.
- [19] A. Hannak, G. Soeller, D. Lazer, A. Mislove, and C. Wilson. Measuring price discrimination and steering on e-commerce web sites. In *Internet Measurement Conference*, pages 305–318, 2014.
- [20] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *NIPS*, pages 3315–3323, 2016.
- [21] W. House. Big data: A report on algorithmic systems, opportunity, and civil rights. *Washington, DC: Executive Office of the President, White House*, 2016.
- [22] D. Koutra, P. N. Bennett, and E. Horvitz. Events and controversies: Influences of a shocking news event on information seeking. In *WWW*, pages 614–624, 2015.
- [23] J. Kulshrestha, M. Eslami, J. Messias, M. B. Zafar, S. Ghosh, I. Shibpur, I. K. P. Gummadi, and K. Karahalios. Quantifying search bias: Investigating sources of bias for political searches in social media. In *CSCW*, 2017.
- [24] B. Lepri, J. Staiano, D. Sangokoya, E. Letouzé, and N. Oliver. The tyranny of data? the bright and dark sides of data-driven decision-making for social good. *arXiv preprint arXiv:1612.00323*, 2016.
- [25] Z. Liu and I. Weber. Is twitter a public sphere for online conflicts? a cross-ideological and cross-hierarchical look. In *SocInfo*, pages 336–347, 2014.
- [26] A. Mowshowitz and A. Kawaguchi. Measuring search engine bias. *Information Processing & Management*, 41(5):1193–1205, 2005.
- [27] L. Muchnik, S. Aral, and S. J. Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146):647–651, 2013.
- [28] B. Pan, H. Hembrooke, T. Joachims, L. Lorigo, G. Gay, and L. Granka. In google we trust: Users decisions on rank, position, and relevance. *Journal of Computer-Mediated Communication*, 12(3):801–823, 2007.
- [29] F. Pasquale. *The black box society: The secret algorithms that control money and information*. Harvard University Press, 2015.
- [30] A. Romei and S. Ruggieri. A multidisciplinary survey on discrimination analysis. *The Knowledge Engineering Review*, 29(05):582–638, 2014.
- [31] C. Sandvig, K. Hamilton, K. Karahalios, and C. Langbort. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry*, 2014.
- [32] J. L. Skeem and C. T. Lowenkamp. Risk, race, and recidivism: predictive bias and disparate impact. *Criminology*, 54(4):680–712, 2016.
- [33] J. Stoyanovich, S. Abiteboul, and G. Miklau. Data, responsibly: Fairness, neutrality and transparency in data analysis. In *EDBT*, 2016.
- [34] L. Sweeney. Discrimination in online ad delivery. *Queue*, 11(3):10, 2013.
- [35] L. Vaughan and M. Thelwall. Search engine coverage bias: evidence and possible causes. *Information processing & management*, 40(4):693–707, 2004.
- [36] I. Weber, V. R. K. Garimella, and A. Batayneh. Secular vs. islamist polarization in egypt on twitter. In *ASONAM*, pages 290–297, 2013.
- [37] F. M. F. Wong, C. W. Tan, S. Sen, and M. Chiang. Quantifying political leaning from tweets and retweets. *ICWSM*, 13:640–649, 2013.
- [38] M. B. Zafar, K. P. Gummadi, and C. Danescu-Niculescu-Mizil. Message impartiality in social media discussions. In *ICWSM*, pages 466–475, 2016.
- [39] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi. Learning fair classifiers. *arXiv preprint arXiv:1507.05259*, 2015.
- [40] M. B. Zafar, I. Valera, M. G. Rodriguez, and K. P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *WWW*, 2017.